MNRAS **000**, 1–6 (2015) Preprint 24 November 2015 Compiled using MNRAS LATEX style file v3.0

Approximating W projection as a separable kernel

Bruce Merry[1]⋆
[1]*SKA South Africa, 3rd Floor, The Park, Park Road, Pinelands, 7405, South Africa*

Accepted 2015 November 20. Received 2015 November 20; in original form 2015 September 10

arXiv:1511.07152v1 [astro-ph.IM] 23 Nov 2015

ABSTRACT
W projection is a commonly-used approach to allow interferometric imaging to be accelerated by Fast Fourier Transforms (FFTs), but it can require a huge amount of storage for convolution kernels. The kernels are not separable, but we show that they can be closely approximated by separable kernels. The error scales with the fourth power of the field of view, and so is small enough to be ignored at mid to high frequencies. We also show that hybrid imaging algorithms combining W projection with either faceting, snapshotting, or W stacking allow the error to be made arbitrarily small, making the approximation suitable even for high-resolution wide-field instruments.

Key words: techniques: interferometric – methods: numerical

1 INTRODUCTION

In interferometric imaging, the relationship between sampled visibilities and the image plane is *almost* a Fourier transform, but with a term that depends on w – the dot product of the baseline vector with the unit vector towards the phase centre. A number of approaches have been developed to deal with this troublesome w term.

One of these is W projection (Cornwell et al. 2008), which converts the image-space multiplication by a phase screen into a convolution with its Fourier transform in uvw space. This convolution is usually combined with an anti-aliasing filter, yielding a combined kernel, the Gridding Convolution Function (GCF), for which closed-form formulae are not known. Instead, the GCF is sampled in three dimensions and stored in a lookup table. Depending on the desired accuracy, these kernels can become excessively large. Apart from requiring large amounts of memory or disk space to store, this can also reduce performance as the samples are moved in and out of caches.

Our contribution is to recognize that the phase screen for a particular value of w, and hence its Fourier transform, is very close to being separable, i.e., an outer product of two one-dimensional functions. By approximating the phase screen in this way, we can drastically reduce the storage required for lookup tables. Conversely, it allows for much finer sampling of the kernel in the same amount of memory, potentially improving accuracy.

Sec. 2 recaps the basics of W projection, and introduces the notation we use. In Sec. 3 we describe our separable approximation, and provide a theoretical analysis of the error. The error in position scales with the fourth power of the field of view, and so is acceptable for small to medium fields of view.

W projection can be combined with other imaging techniques, which we discuss in Sec. 4. We show that this allows the error to be reduced further, allowing our approach to be used even for wide fields of view. Sec. 5 discusses the effect on computation cost, and our conclusions are presented in Sec. 6.

2 BACKGROUND AND NOTATION

We will focus mainly on transforming from sampled visibilities to a dirty image. Nevertheless, the same techniques and analysis apply to prediction of visibilities from a model image.

As usual, u, v, w are baseline coordinates in a fixed coordinate system with w in the direction of the phase centre. We measure u, v, w in wavelengths rather than units of distance. The corresponding direction cosines are denoted l, m, n, with $n = \sqrt{1 - l^2 - m^2}$. For compactness of notation, we will also use $\boldsymbol{l} = \begin{pmatrix} l & m \end{pmatrix}^T$ and $\boldsymbol{u} = \begin{pmatrix} u & v \end{pmatrix}^T$ interchangeably.

We will ignore direction-dependent effects, and assume a time- and baseline-independent perceived brightness distribution $I(\boldsymbol{l})$. Let the ith visibility have coordinates (u_i, v_i, w_i), visibility value V_i and weight W_i. The dirty image I^{D} is given by

$$\frac{I^{\mathrm{D}}(\boldsymbol{l})}{n} = \sum_i W_i V_i \mathrm{e}^{2\pi\mathrm{i}(\boldsymbol{u}_i \cdot \boldsymbol{l} + w_i(n-1))}. \quad (1)$$

The corresponding prediction of visibilities from a model I is

$$V_i = \iint \frac{I(\boldsymbol{l})}{n} \mathrm{e}^{-2\pi\mathrm{i}(\boldsymbol{u}_i \cdot \boldsymbol{l} + w_i(n-1))} \mathrm{d}\boldsymbol{l}. \quad (2)$$

⋆ E-mail: bmerry@ska.ac.za

© 2015 The Authors

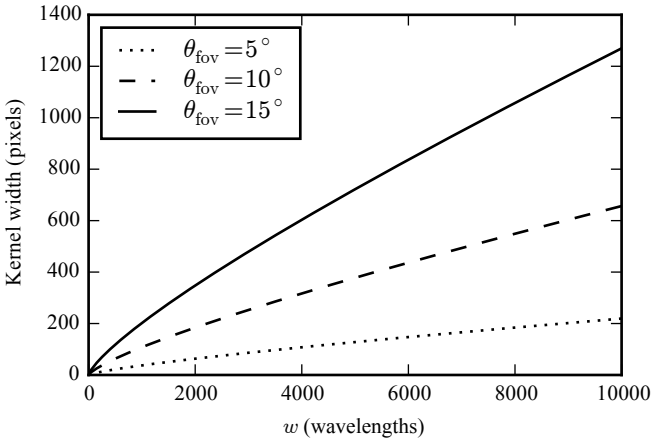

Figure 1. Support for the W kernel with truncation at $\eta = 0.01$ of the peak. θ_{fov} is the field of view.

Let $g_w(\boldsymbol{l}) = \mathrm{e}^{2\pi \mathrm{i} w(n-1)}$ and $G_w(\boldsymbol{u})$ be its Fourier transform, and let $V_i^{\mathrm{W}}(u,v) = \delta(u-u_i)\delta(v-v_i)W_iV_i$. Then

$$\begin{aligned} I^{\mathrm{D}}(l,m)/n &= \sum_i \left(\iint V_i^{\mathrm{W}}(\boldsymbol{u})\,\mathrm{d}\boldsymbol{u} \right) \mathrm{e}^{2\pi\mathrm{i}(\boldsymbol{u}_i\cdot\boldsymbol{l}+w_i(n-1))} \\ &= \sum_i \left(\iint V_i^{\mathrm{W}}(\boldsymbol{u}) \mathrm{e}^{2\pi\mathrm{i}\boldsymbol{u}\cdot\boldsymbol{l}}\,\mathrm{d}\boldsymbol{u} \right) g_{w_i}(\boldsymbol{l}) \\ &= \sum_i \mathsf{F}^{-1}[V_i^{\mathrm{W}}](\boldsymbol{l}) g_{w_i}(\boldsymbol{l}) \\ &= \sum_i \mathsf{F}^{-1}[V_i^{\mathrm{W}} * G_{w_i}](\boldsymbol{l}) \\ &= \mathsf{F}^{-1}\Big[\sum_i V_i^{\mathrm{W}} * G_{w_i} \Big](\boldsymbol{l}). \end{aligned} \tag{3}$$

This is typically combined with an antialiasing kernel C (with inverse Fourier transform c) to give

$$I^{\mathrm{D}}(l,m)/n = \mathsf{F}^{-1}\Big[\sum_i V_i^{\mathrm{W}} * G_{w_i} * C \Big](l,m)/c(l,m). \tag{4}$$

The two convolution kernels G_w and C are combined into a gridding convolution function (GCF) F_w by multiplying their inverse Fourier transforms g_w and c, and taking the Fourier transform of the result.

Generating a kernel is expensive, so usually they are either fully precomputed, or are generated on-demand then cached. The storage size is thus an important consideration. The antialiasing kernel C is typically small (e.g., 9×9 pixels), so the support of the GCF is dominated by the support of G_w. The support (in uv-plane pixels) necessary to represent the function out to the fraction η of the peak is (Mitchell & Bernardi 2014)

$$2\theta_{\mathrm{fov}} \sqrt{\left(\frac{w\theta_{\mathrm{fov}}}{2}\right)^2 + \frac{w^{3/2}\theta_{\mathrm{fov}}}{2\pi\eta}}. \tag{5}$$

Fig. 1 shows typical values of this function for $\eta = 0.01$.

Kernels can be hundreds of pixels across, even for baselines of a few kilometres. What is worse, they must be oversampled to avoid aliasing effects. A typical oversampling factor is 8 (Romein 2012), which means that a single kernel may require millions of samples and thus consume tens of megabytes.

Cornwell et al. (2012) suggest that the number of w planes should be

$$\frac{\pi w_{\max}(\theta_{\mathrm{fov}})^2}{\sqrt{2\Delta A}} \tag{6}$$

where ΔA is the tolerable loss of amplitude in the image plane due to decorrelation effects. This can easily reach 10^4 at low elevations with $\Delta A = 0.01$, and storing all the kernels can require tens or even hundreds of gigabytes. Our method can reduce memory usage by a factor of 1000 or more, making precomputation practical.

3 DERIVATION AND ANALYSIS

We aim to approximate F_w as a separable function. We assume that the antialiasing function c is separable, and hence it suffices to approximate g_w as a separable function.

Recall that $g_w(l,m) = \mathrm{e}^{2\pi\mathrm{i}w(n-1)}$ and $n = \sqrt{1-l^2-m^2}$. Using the Taylor expansion $\sqrt{1-x} = 1 - \frac{1}{2}x - \frac{1}{8}x^2 + O(x^3)$, we get

$$g_w(l,m) = \mathrm{e}^{2\pi\mathrm{i}w\left[-\frac{1}{2}(l^2+m^2)-\frac{1}{8}(l^2+m^2)^2+O(l^6+m^6)\right]}. \tag{7}$$

The first few terms of the phase depend on either l or m, but not both, and these are the dominant terms when $|l|, |m| \ll 1$. This is what makes g approximately separable. Let

$$g_w^1(l) = \mathrm{e}^{2\pi\mathrm{i}w\left[-\frac{1}{2}l^2-\gamma l^4\right]} \tag{8}$$

$$\widetilde{g}_w(l,m) = g_w^1(l) g_w^1(m), \tag{9}$$

where γ is a tuning parameter we will discuss later. We approximate g_w by $\widetilde{g}_w$. This approximation introduces an image-space phase error of

$$\begin{aligned} \Delta\phi(\boldsymbol{l}, w) &= 2\pi w\Big[-(\tfrac{1}{2}l^2 + \gamma l^4 + \tfrac{1}{2}m^2 + \gamma m^4) \\ &\qquad -(\sqrt{1-l^2-m^2}-1)\Big] \\ &= 2\pi w\Big[\tfrac{1}{2}(l^2+m^2) + \tfrac{1}{8}(l^2+m^2)^2 + O(l^6+m^6) \\ &\qquad -\tfrac{1}{2}l^2 - \gamma l^4 - \tfrac{1}{2}m^2 - \gamma m^4 \Big] \\ &= 2\pi w\Big[(\tfrac{1}{8}-\gamma)(l^4+m^4) + \tfrac{1}{4}l^2m^2 + O(l^6+m^6) \Big]. \end{aligned} \tag{10}$$

Denote the factor inside the square brackets by $\epsilon(\boldsymbol{l})$. We would like to minimize $|\epsilon(\boldsymbol{l})|$ over the image. The largest values will clearly be along the edges. Fig. 2 shows how $\epsilon(\boldsymbol{l})$ varies along an edge for several values of γ. Ignoring the $O(l^6+m^6)$ term, the value $\gamma = \frac{5}{24}$ minimizes the maximum error at $\frac{1}{12}m_{\max}^4$, while $\gamma = \frac{1}{8}$ favors accuracy on the $l = 0$ and $m = 0$ axes and pushes the error into the corners.

We will now approximate the array as instantaneously co-planar (ignoring the slight curvature due to the shape of the Earth), so that w can be computed as $w = \boldsymbol{d}\cdot\boldsymbol{u}$ for some time-varying $\boldsymbol{d}$. Note that $\|\boldsymbol{d}\| = \cot a$ (a being the elevation angle), so it is bounded by the minimum elevation. This allows us to recast the phase error as an instantaneous position error. When substituting the approximation into (3), we get

$$\begin{aligned} I^{\mathrm{D}}(l,m)/n &= \sum_i W_iV_i \mathrm{e}^{2\pi\mathrm{i}\left[\boldsymbol{u}_i\cdot\boldsymbol{l}+w_i\left(\sqrt{1-\|\boldsymbol{l}\|^2}-1+\epsilon(\boldsymbol{l})\right)\right]} \\ &= \sum_i W_iV_i \mathrm{e}^{2\pi\mathrm{i}\boldsymbol{u}_i\cdot\left[\boldsymbol{l}+\left(\sqrt{1-\|\boldsymbol{l}\|^2}-1+\epsilon(\boldsymbol{l})\right)\boldsymbol{d}\right]}. \end{aligned} \tag{11}$$

If we can find $\boldsymbol{l}'$ such that

$$\boldsymbol{l}' + (\sqrt{1-\|\boldsymbol{l}'\|^2}-1)\boldsymbol{d} = \boldsymbol{l} + (\sqrt{1-\|\boldsymbol{l}\|^2}-1+\epsilon(\boldsymbol{l}))\boldsymbol{d} \tag{12}$$

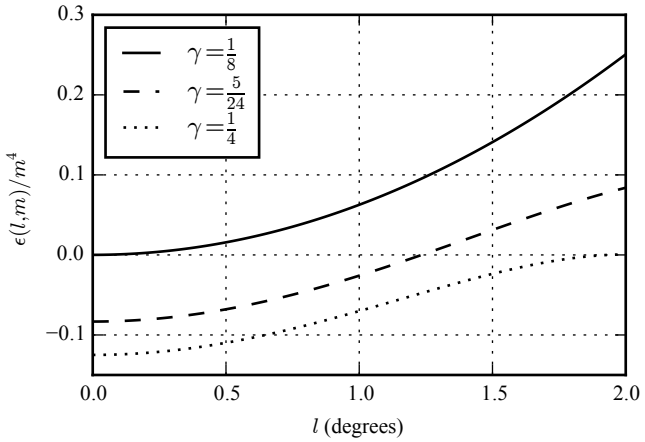

Figure 2. Phase error factor along the edge of an image with 4° field of view ($m = 2°$).

then the phase error is equivalent to shifting $\boldsymbol{l}'$ in the correct dirty image to $\boldsymbol{l}$ in the approximated dirty image. In other words, at a single point in time the approximation will cause a distortion in the image; as the Earth rotates, $\boldsymbol{d}$ will change, causing sources to be smeared out.

If $\Delta\boldsymbol{l} = \boldsymbol{l}' - \boldsymbol{l}$ is sufficiently small, then this error can reasonably be ignored. Cornwell et al. (2012) show that the relative loss in amplitude at $\boldsymbol{l}$ is proportional to $\left(\frac{\|\Delta\boldsymbol{l}\|}{\theta_{\text{res}}}\right)^2$, assuming a parabolic dirty beam, where θ_{res} is the resolution; it is thus desirable that $\|\Delta\boldsymbol{l}\| \ll \theta_{\text{res}}$. While (12) can be solved exactly (see Appendix A), we can obtain more insight from the approximation

$$\Delta\boldsymbol{l} \approx (1 + \boldsymbol{l}\cdot\boldsymbol{d}/n)\epsilon(\boldsymbol{l})\boldsymbol{d} \approx \epsilon(\boldsymbol{l})\boldsymbol{d} \tag{13}$$

which suggests that $\|\epsilon(\boldsymbol{l})\boldsymbol{d}\| \ll \theta_{\text{res}}$ is sufficient.

We can translate this into a relationship between parameters of the instrument. The field of view θ_{fov} is given by $2Q_{\text{fov}}\frac{\lambda}{D}$, where λ is the wavelength, D is the antenna/station diameter, and Q_{fov} is a constant that depends on the antenna (e.g., illumination tapering) and the desired field of view relative to the primary beam size. For example, $Q_{\text{fov}} = 1.22$ gives an image that encloses the first null of an ideal Airy disk (Taylor et al. 1999, p. 41). For resolution we use the rule of thumb $\theta_{\text{res}} = \frac{3\lambda}{2B}$, where B is the length of the longest baseline; the actual constant factor depends on the *uv* coverage and imaging weights (Taylor et al. 1999, p. 131). As noted above, $\|\boldsymbol{d}\| = \cot a$. Thus,

$$\frac{\|\Delta\boldsymbol{l}\|}{\theta_{\text{res}}} \approx \frac{\frac{1}{12}(Q_{\text{fov}}\frac{\lambda}{D})^4\cot a}{\left(\frac{3\lambda}{2B}\right)} = \frac{1}{18}Q_{\text{fov}}^4\lambda^3 D^{-4}B\cot a \tag{14}$$

To provide an example, let us consider some worst-case values for MeerKAT (Karoo Array Telescope) (SKA South Africa 2015, 2014): $\lambda = 0.333$ m (900 MHz), $D = 13.5$ m, $B = 8$ km, $a = 15°$ gives $\frac{\|\Delta\boldsymbol{l}\|}{\theta_{\text{res}}} \approx 0.0018Q_{\text{fov}}^4$, allowing for a large Q_{fov} without problems.

Phase 1 of Square Kilometre Array mid-frequency instrument (SKA1-MID) will support lower frequencies and much longer baselines (SKA Organisation HQ 2015): $D = 15$ m, $B = 150$ km, $\lambda = 0.857$ m (350 MHz). In this case, (14) gives $0.39Q_{\text{fov}}^4$, assuming the same 15° minimum elevation as for MeerKAT. This means that the position error is of the same order as the beam size, which over time will cause decoherence at the edges of the image.

Due to the cubic dependence on λ, our basic approach is applicable at mid and high frequencies, but breaks down at lower frequencies. In general, we see limited value to the basic approach below about 500 MHz, at least for arrays of small dishes. However, combining W projection with other imaging approaches allows the error to be made arbitrarily small, making our approach suitable even at low frequencies. This is discussed further in Sec. 4.

It is interesting to note that the analysis above applies to other phase errors of the form $2\pi w\epsilon(\boldsymbol{l})$. In particular, if no correction for W is done at all, then $\epsilon(\boldsymbol{l}) = 1 - \sqrt{1 - \|\boldsymbol{l}\|^2}$. Surprisingly, (13) does not depend on the array diameter or largest w value, even though w effects are conventionally associated with long baselines (Cornwell et al. 2008). Long baselines play an indirect role, in that they give high resolution and hence position errors are larger relative to the synthesized beam.

4 COMBINATION WITH OTHER TECHNIQUES

W projection is only one approach to dealing with the w term, and hence allowing Fast Fourier Transforms to be used to accelerate imaging. Others include faceting, snapshot imaging, and W stacking. In each case, it is possible to hybridize the algorithm with W projection to combine the computational strengths of both. W projection's strength is in handling small w values very accurately (limited only by sampling resolution), but it becomes computationally very costly to handle large w values as the necessary support of the GCF grows. Hybrid techniques use alternative methods to make coarse phase corrections, with the finer corrections left to W projection.

Since the phase error in our method scales with w, it seems likely that these hybrid methods would be particularly suitable since the w values involved would be reduced. The following subsections show that this is indeed the case.

4.1 Faceting

Faceting has a number of advantages for imaging: each facet requires less working memory, and the reduction in field of view allows w effects to be ignored or to be corrected more easily. Here we will discuss faceting in which the field of view of each facet is small, but not so small that w effects can be completely ignored. This is an important domain, because there is a fixed cost per facet to adjust and grid the visibilities, and hence a large number of tiny facets can be computationally costly.

In classical faceting, the coordinate system is rotated and the phase of visibilities is adjusted so that each facet has its own phase centre. In terms of w effects, each facet behaves as an independent image with a narrower field of view. Since the error scales with the fourth power of field of view, using even a small number of facets will greatly reduce the errors.

More recently, 'image-plane' faceting (Kogan & Greisen 2009) applies a non-orthogonal change of coordinate system so that the facets are all part of a single plane tangent to the celestial sphere at the original phase centre, rather than

each being tangent at the facet centre. This simplifies generation of a single image from the facets, but complicates our analysis.

Kogan & Greisen (2009) derive the equations for image-plane faceting using only a first-order Taylor approximation to the w factor. In their derivation, all w effects are corrected by suitable adjustments to $\boldsymbol{u}$. For a hybrid W-projection/faceting algorithm, we must compute a second-order approximation.

Let $h(\boldsymbol{x}) = \sqrt{1 - \|\boldsymbol{x}\|^2}$, let $\boldsymbol{l}_0$ be the facet centre, with $n_0 = h(\boldsymbol{l}_0)$, and let $\boldsymbol{l}$ refer to the position relative to the *facet* centre, such that the position relative to the original phase centre is $\boldsymbol{l} + \boldsymbol{l}_0$. The modified w correction is

$$g_w(\boldsymbol{l}) = \mathrm{e}^{2\pi \mathrm{i} w(n-n_0)} = \mathrm{e}^{2\pi \mathrm{i} w[h(\boldsymbol{l}_0+\boldsymbol{l})-h(\boldsymbol{l}_0)]} \tag{15}$$

We now approximate h using a Taylor polynomial about $\boldsymbol{l}_0$:

$$\begin{aligned} h(\boldsymbol{l}_0 + \boldsymbol{l}) - h(\boldsymbol{l}_0) &\approx \nabla h(\boldsymbol{l}_0)\boldsymbol{l} + \tfrac{1}{2}\boldsymbol{l}^T \mathbf{H} h(\boldsymbol{l}_0)\boldsymbol{l} \\ &= -n_0^{-1}\boldsymbol{l}_0^T\boldsymbol{l} + \tfrac{1}{2}\boldsymbol{l}^T[-n_0^{-1}I - n_0^{-3}\boldsymbol{l}_0\boldsymbol{l}_0^T]\boldsymbol{l} \\ &= -n_0^{-1}\boldsymbol{l}_0^T\boldsymbol{l} - \tfrac{1}{2}(n_0^{-1} + l_0^2 n_0^{-3})l^2 \\ &\quad - \tfrac{1}{2}(n_0^{-1} + m_0^2 n_0^{-3})m^2 - n_0^{-3} l_0 m_0 lm. \end{aligned} \tag{16}$$

The initial term is absorbed into the coordinate transformation $\boldsymbol{u}' = \boldsymbol{u} - n_0^{-1}\boldsymbol{l}_0 w$. The next two terms can be handled by a separable kernel, while the final term is not separable. We replace (9) by

$$g_w^{\mathrm{l}}(l) = \mathrm{e}^{2\pi \mathrm{i} w\left[-\frac{1}{2}(n_0^{-1} + l_0^2 n_0^{-3})l^2\right]} \tag{17}$$

$$g_w^{\mathrm{m}}(m) = \mathrm{e}^{2\pi \mathrm{i} w\left[-\frac{1}{2}(n_0^{-1} + m_0^2 n_0^{-3})m^2\right]} \tag{18}$$

$$\widetilde{g}_w(l, m) = g_w^{\mathrm{l}}(l) g_w^{\mathrm{m}}(m). \tag{19}$$

There are a few noteworthy differences from the case of classical faceting. First, the component functions are now facet-dependent, and different for l and m. We have also used a lower-order approximation, because without being able to incorporate the $-n_0^{-3} l_0 m_0 lm$ term, there is little point in trying to compute higher-order terms. This also means that for a given total field of view, the error now only scales with the square rather than the fourth power of the facet width.

Fig. 3 shows the position error at 15° elevation and a 10° field of view, using 11×11 and 5×5 grids of facets. In the 11×11 case, the error is reduced by an order of magnitude compared to the unfaceted approach. However, the discontinuities at the boundaries between facets may be problematic, since a source on a boundary will appear in different positions in the two facets.

Our claim that the error scales quadratically with facet size holds between 3a and 3b, but 3c does not follow the same scaling. This is for two reasons. Firstly, 3c uses (9) rather than (19), and the former partially corrects for fourth-order errors. Secondly, the claim only holds for facets with the same centre; since 3c has the facet centre at the phase centre, the errors are smaller.

4.2 W snapshots

Snapshot imaging corrects the position error introduced by w terms in image space. Since the position error is time-varying, an image is created for a small window of time, then reprojected to correct the distortions. Over a longer observation, many images are made and added together. For each snapshot, a single $\boldsymbol{d}_0$ value is found such that $w \approx \boldsymbol{u} \cdot \boldsymbol{d}_0$, and then after creating the image, the image is reprojected so that $\boldsymbol{l}^{\mathrm{s}} = \boldsymbol{l} + \boldsymbol{d}_0(n-1)$ is mapped to $\boldsymbol{l}$.

W snapshots (Cornwell et al. 2012) is a hybrid method combining snapshot imaging with W projection. The snapshot warping corrects for an average position error over the snapshot duration, while W projection fixes up residual errors caused by long snapshots or by non-planar arrays. This allows the snapshot duration to be much longer than for traditional snapshot imaging, while still reducing the gridding costs by keeping w values small.

The component $2\pi \mathrm{i} \boldsymbol{u} \cdot \boldsymbol{d}_0(n-1)$ of the phase is absorbed into the reprojection, leaving a residual of $2\pi \mathrm{i} \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)(n-1)$ to be handled by W projection[1]. It follows that the remaining position error from using a separable kernel is

$$\Delta \boldsymbol{l} \approx (1 + \boldsymbol{l} \cdot (\boldsymbol{d} - \boldsymbol{d}_0))\epsilon(\boldsymbol{l})(\boldsymbol{d} - \boldsymbol{d}_0). \tag{20}$$

There is another source of phase error, not discussed by Cornwell et al. (2012): W projection does not account for the reprojection[2], and hence the actual phase correction applied is $2\pi \mathrm{i} \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)(\sqrt{1 - \|\boldsymbol{l}^{\mathrm{s}}\|^2} - 1)$. This will introduce a phase error of

$$\begin{aligned} &2\pi \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)\left(\sqrt{1 - \|\boldsymbol{l}^{\mathrm{s}}\|^2} - \sqrt{1 - \|\boldsymbol{l}\|^2}\right) \\ &\approx 2\pi \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)(-\boldsymbol{l} \cdot \boldsymbol{d}_0(n-1)) \\ &\approx 2\pi \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)(\tfrac{1}{2}\|\boldsymbol{l}\|^2 \boldsymbol{l} \cdot \boldsymbol{d}_0) \end{aligned} \tag{21}$$

For comparison, the phase error associated with separation of the kernel, when $m = 0$ and $\gamma = \frac{5}{24}$, is approximately $2\pi \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)\epsilon(\boldsymbol{l}) \approx 2\pi \boldsymbol{u} \cdot (\boldsymbol{d} - \boldsymbol{d}_0)(\frac{1}{12}\|\boldsymbol{l}\|^4)$. Unless the phase centre is close to the zenith, $\|\boldsymbol{d}_0\| \gg \|\boldsymbol{l}\|$, and so the error due to separation will be insignificant compared to this other error. Of course, both errors can be made arbitrarily small by limiting the snapshot length, although in the limit the assumption of a coplanar array will break down.

4.3 W stacking

W stacking (Offringa et al. 2014) is similar to snapshot imaging in that visibilities are partitioned, and within each partition the average w effects are corrected in image space. Instead of partitioning visibilities by time, they are partitioned by w, and the image-space adjustment is a phase correction rather than a reprojection. As with snapshotting, it is possible to use coarser partitioning, and hence fewer Fourier transforms, by using W projection to deal with residual w terms.

Unlike in the previous cases, the w projection term is no longer directly correlated with $\boldsymbol{u}$ at a point in time. The errors will thus cause smearing rather than a systematic shift in position. Let Δw be the residual w that will be adjusted by w projection. If the slices are evenly spaced $2w_{\max}$ apart, then $|\Delta w| \le w_{\max}$. With sufficiently many slices, Δw can be treated as a uniformly-distributed random variable.

A phase error of ϕ will cause a relative amplitude loss

[1] For analysis we will assume a coplanar array, but deviations from planarity will also appear in the residual.

[2] This could be fixed by generating a custom kernel for each value of $\boldsymbol{d}_0$. We have not analyzed this approach.

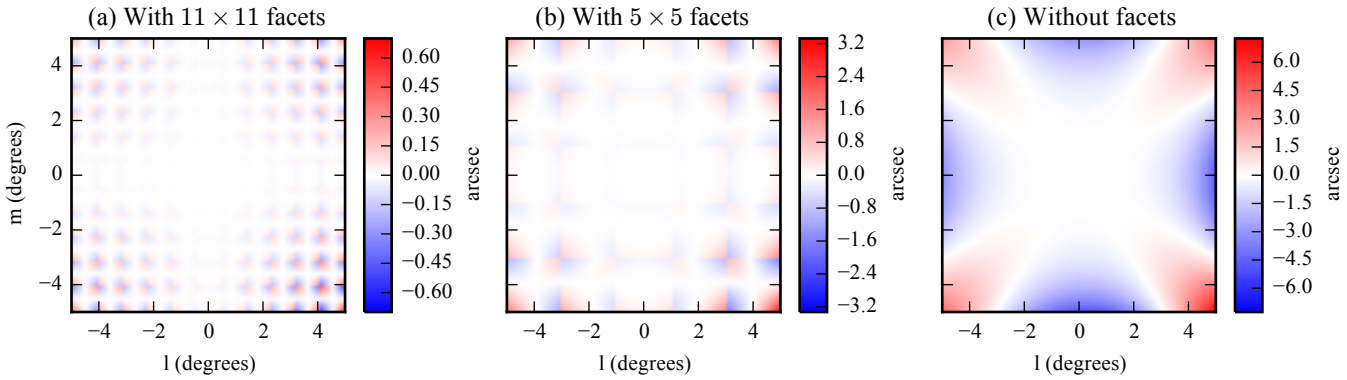

Figure 3. Position error when using (19) with facets (a and b) versus (9) and no facets (c).

of $1 - \cos\phi$ at the position of a point source. Hence, the expected loss of amplitude is

$$\begin{aligned}\langle 1 - \cos 2\pi\Delta w\epsilon(\boldsymbol{l})\rangle &\approx \langle \tfrac{1}{2}(2\pi\Delta w\epsilon(\boldsymbol{l}))^2\rangle \\ &= 2\pi^2\epsilon(\boldsymbol{l})^2\langle\Delta w^2\rangle \\ &= \tfrac{2}{3}\pi^2\epsilon(\boldsymbol{l})^2 w_{\max}^2.\end{aligned} \tag{22}$$

Let us re-evaluate the SKA1-MID example from Sec. 3, with a 10° field of view (corresponding very roughly to the second null of the primary beam). In this case, $|\epsilon(\boldsymbol{l})| < 10^{-5}$ everywhere in the image, so provided that $w_{\max} \ll 10^5$, the loss in amplitude will be minimal. At 500 MHz, the longest baseline is 2.5×10^5 wavelengths, so this will not require an unreasonable number of slices.

A potential issue with a core-heavy array is that if the $w = 0$ slice is too wide, it will contain a substantial fraction of the visibilities. These visibilities contribute to the dirty image exactly as if no W stacking was applied, and hence create a ghost source at $\boldsymbol{l}'$ as in equation (12). This suggests placing a large number of slices close to $w = 0$, with wider-spaced slices at larger w values. Adapting slice width to the visibility density will also reduce computation time, because narrow slices need less support for the convolution kernel, hence speeding up gridding, but sparse regions benefit from fewer, wider slices to reduce Fourier transform costs.

5 COMPUTATION COST

With our approach, the gridding convolution function becomes a separable function $\widetilde{F}_w(u,v) = F_w^1(u)F_w^1(v)$. Updating the grid cell at (u,v) with visibility i now requires computing $V_i F_{w_i}^1(u-u_i)F_{w_i}^1(v-v_i)$ rather than $V_i F_{w_i}(u-u_i, v-v_i)$. This would appear to require double the number of multiplications. However, the partial product $V_i F_{w_i}^1(u-u_i)$ depends only on u and not v, and hence can be computed once and reused for every value of v in the support of the kernel. The number of multiplications thus only increases by a factor of $1 + \frac{1}{|F_w^1|}$, where $|F_w^1|$ is the width of the kernel in grid cells.

It is also important to emphasize that floating-point operations are not necessarily the bottleneck in gridding: Romein (2012) and Muscat (2014) both report that performance is limited by the memory system in their respective GPU-accelerated gridders. Muscat reports that loading the kernel data from texture memory is the limiting factor, and our smaller kernel is likely to improve cache efficiency here. Our prototype GPU gridder using a separable kernel achieves greater efficiency than the published figures for either of these implementations, but further work is required to establish whether this is due to other factors.

6 CONCLUSIONS AND FUTURE WORK

We have demonstrated a method to approximate the GCF for W projection which reduces the lookup table from three to two dimensions, leading to a massive reduction in memory. This makes it possible to store more w planes, to make better usage of caches, and possibly to keep the entire kernel in a smaller but faster level of a memory hierarchy. We have focused on dirty imaging, but the same approximation can be used for prediction, allowing for Cotton–Schwab major cycles (Schwab 1984).

The approximation introduces a phase error, which is linear in w and hence translates to an instantaneous position error for a coplanar array. The error is small, scaling with the fourth power of the field of view. The approach is thus practical for mid- and high-frequency dish arrays, even with pure W projection, provided that baselines are no more than a few tens of kilometres long. The error can also be made arbitrarily small by using faceting, snapshotting or W stacking with a sufficient number of facets, snapshots or stack slices. This allows the basic idea to be extended to higher resolutions and lower frequencies, such as is envisioned for SKA1-MID. It remains future work to investigate whether wide-field, low-frequency aperture arrays can be supported without requiring an excessively large number of facets/snapshots/slices that would harm the overall performance.

While we believe that our approach will increase performance due to reduced pressure on memory systems, we have not yet implemented both versions of the kernel in a single gridding application to provide a fair comparison.

An idea that merits further investigation is to obtain better accuracy by storing and applying the correction $F_w - \widetilde{F}_w$. This will potentially require far less support and/or sampling rate than the original GCF, and hence can be stored in a small yet full-dimension table. It may also be possible to approximate this correction as another separable function, but rotated 45° to the original.

Another method used in interferometric imaging is A projection, in which primary beam correction is done in the

uv plane by convolution (Bhatnagar et al. 2006). This may also be combined with W projection, and one may naturally ask whether our technique can be applied to the combination. Unfortunately this is unlikely to be as successful, as the primary beam rotates on the sky and is unlikely to be separable in all orientations. Our technique can still be used with other approaches to primary beam correction, such as image-space correction for each snapshot with W snapshots.

REFERENCES

Bhatnagar S., Golap K., Cornwell T., 2006, Correction of Errors due to Antenna Power Patterns during Imaging, EVLA memo 100

Cornwell T., Golap K., Bhatnagar S., 2008, IEEE J. of Selected Topics in Signal Processing, 2, 647

Cornwell T. J., Voronkov M. A., Humphreys B., 2012, Proc. SPIE, 8500, 85000L

Kogan L., Greisen E. W., 2009, Faceted imaging in AIPS, AIPS memo 113

Mitchell D., Bernardi G., 2014, Analysis of w-projection kernel size, SKA SDP Consortium document SKA-TEL-SDP-IMG-Wkernels

Muscat D., 2014, Master's thesis, University of Malta

Offringa A. R., et al., 2014, MNRAS, 444, 606

Romein J. W., 2012, in Proc. 26th ACM Int. Conf. on Supercomputing. ICS '12. ACM, New York, NY, USA, pp 321–330

SKA Organisation HQ 2015, The world's largest radio telescope takes a major step towards construction, `https://www.skatelescope.org/news/worlds-largest-radio-telescope-near-construction/`

SKA South Africa 2014, South Africa's MeerKAT Radio Telescope, `https://docs.google.com/a/ska.ac.za/viewer?a=v&pid=sites&srcid=c2thLmFjLnphfHB1YmxpY3xneDo3Yjg5ODU3ZGYxMjI4NDg0`

SKA South Africa 2015, MeerKAT Array Releases and Specifications, `http://public.ska.ac.za/meerkat/schedule`

Schwab F. R., 1984, Astron. J., 89, 1076

Taylor G. B., Carilli C. L., Perley R. A., eds, 1999, Synthesis Imaging in Radio Astronomy II Astronomical Society of the Pacific Conference Series Vol. 180

APPENDIX A: PROOF OF APPROXIMATION FOR $\Delta\boldsymbol{l}$

Here we will derive an exact solution to (12) and justify the approximation given in (13). Rearranging (12) gives

$$\Delta\boldsymbol{l} = \left[\sqrt{1-\|\boldsymbol{l}\|^2} - \sqrt{1-\|\boldsymbol{l}+\Delta\boldsymbol{l}\|^2} + \epsilon(\boldsymbol{l})\right]\boldsymbol{d}. \tag{A1}$$

It is immediately clear that $\Delta\boldsymbol{l}$ is a multiple of $\boldsymbol{d}$, so let $\Delta\boldsymbol{l} = y\boldsymbol{d}$. Let

$$k = \sqrt{1-\|\boldsymbol{l}\|^2} + \epsilon(\boldsymbol{l}) = n + \epsilon(\boldsymbol{l}). \tag{A2}$$

Then

$$\begin{aligned}
& y = \sqrt{1-\|\boldsymbol{l}\|^2} - \sqrt{1-\|\boldsymbol{l}+y\boldsymbol{d}\|^2} + \epsilon(\boldsymbol{l}) && \text{(A3)}\\
\iff\quad & y-k = -\sqrt{1-\|\boldsymbol{l}+y\boldsymbol{d}\|^2} && \text{(A4)}\\
\implies\quad & (y-k)^2 = 1 - \|\boldsymbol{d}\|^2y^2 - 2(\boldsymbol{l}\cdot\boldsymbol{d})y - \|\boldsymbol{l}\|^2 && \text{(A5)}\\
\iff\quad & 0 = (1+\|\boldsymbol{d}\|^2)y^2 + 2(\boldsymbol{l}\cdot\boldsymbol{d}-k)y && \text{(A6)}\\
& \qquad + (k^2 + \|\boldsymbol{l}\|^2 - 1) &&\\
\iff\quad & 0 = (1+\|\boldsymbol{d}\|^2)y^2 + 2(\boldsymbol{l}\cdot\boldsymbol{d}-k)y && \text{(A7)}\\
& \qquad + (2n+\epsilon(\boldsymbol{l}))\epsilon(\boldsymbol{l}). &&
\end{aligned}$$

This is a quadratic equation in y, which can easily be solved. It is possible that both two roots satisfy (A4), but we are only interested in the one with the smaller absolute value.

We now turn to approximating y. To get an initial estimate, we ignore the term $\sqrt{1-\|\boldsymbol{l}\|^2} - \sqrt{1-\|\boldsymbol{l}+\Delta\boldsymbol{l}\|^2}$, giving $y_0 = \epsilon(\boldsymbol{l})$. To improve this, we note that

$$\begin{aligned}
\sqrt{1-\|\boldsymbol{l}+\Delta\boldsymbol{l}\|^2} &= \sqrt{1-\|\boldsymbol{l}+y\boldsymbol{d}\|^2}\\
&= \sqrt{1-(\|\boldsymbol{d}\|^2y^2 + 2(\boldsymbol{l}\cdot\boldsymbol{d})y + \|\boldsymbol{l}\|^2)}\\
&\approx \sqrt{1-\|\boldsymbol{l}\|^2} - \left(\boldsymbol{l}\cdot\boldsymbol{d}/\sqrt{1-\|\boldsymbol{l}\|^2}\right)y\\
&= n - (\boldsymbol{l}\cdot\boldsymbol{d}/n)y
\end{aligned} \tag{A8}$$

where the approximation is a linear Taylor polynomial around $y = 0$. Substituting this back into (A1) gives

$$y \approx (\boldsymbol{l}\cdot\boldsymbol{d}/n)y + \epsilon(\boldsymbol{l}) \tag{A9}$$

and if we use $y_0 = \epsilon(\boldsymbol{l})$ in place of y in the right-hand side, we get the improved approximation

$$y_1 = (1 + \boldsymbol{l}\cdot\boldsymbol{d}/n)\epsilon(\boldsymbol{l}). \tag{A10}$$

This paper has been typeset from a TEX/LATEX file prepared by the author.